

A6826A – PCI-X Dual Channel 2Gb/s Fibre Channel Adapter

Performance Paper for Integrity Servers



Table of contents

Introduction	2
Executive summary	2
Test results	3
IOPs	3
Service demand	3
Throughput	4
Scalability	5
System configuration guidelines	7
Test details	8
Products used in testing	8
Test configuration	9
Additional information	9

Introduction

This paper provides basic I/O performance and scalability information for the A6826A dual-channel 2Gb Fibre Channel adapter on HP's Itanium-based Integrity servers. A series of tests were conducted to evaluate the adapter's performance on rx2600, rx4640, rx7620, rx8620 and HP Integrity Superdome servers. The tests measured the IOPs and throughput of single and multiple adapters.

This paper addresses the performance capabilities of the A6826A adapter when used with the above mentioned HP Integrity servers. This paper also discusses system setup considerations necessary to obtain the maximum possible performance from A6826A adapters installed in these servers.

This paper focuses on the following topics:

- **Test results:** Single and multiple adapter performance data, including IOPs and throughput for a single port and dual ports, will be discussed.
- **Scalability:** Multiple adapter scalability tests on rx2600, rx4640, rx7620, rx8620 and Integrity Superdome systems will be discussed.
- **System configuration guidelines:** rx2600, rx4640, rx7620, rx8620 and Integrity Superdome system configurations and recommendations will be discussed.
- **Test details:** The HP products used, the test setup, the benchmark tool used, and the system configuration in the test setup, will be discussed.

Executive summary

A6826A adapters installed in HP Itanium-based Integrity servers offer an excellent SAN solution. The A6826A adapter provides read throughput of 195MB/s, and write throughput of 186MB/s on a single port running at 2Gb/s. With both ports running at 2Gb/s, a read throughput of 390MB/s, and a write throughput of 372MB/s is achieved.

A6826A performance summary							
Operation		single port			dual ports		
		read	write	bd*	read	write	bd*
IOPs		35180	20619	26187	70032	41221	52414
throughput (MB/s)	achieved	195	186	370	390	372	720
	% of 2Gb/s FC theoretical max	98	93	92.5	97.5	93	90

NOTE:

*bd: Bidirectional Operation

The above performance data was obtained by installing an A6826A adapter in a PCI-X 133MHz slot (Slot 8) of an rx4640 server.

An A6826A adapter installed in a PCI-X slot of an rx4640 server offers excellent bi-directional performance. A single A6826A port running at 2Gb/s provides bi-directional throughput of 370MB/s and both A6826A ports running at 2Gb/s provide bi-directional throughput of 720MB/s with linear scaling.

On other Integrity servers the A6826A adapter provides outstanding performance with linear scaling for up to 4 adapters when all eight ports are being used at 2Gb/s bandwidth. Additional adapters may be added to provide greater connectivity.

Test results

HP obtained the performance data presented in this paper using the diskbench (db) benchmarking utility. Various block size tests were executed for read and write on single and dual ports. The IOPs metric was obtained with 1 KB block size transfers. The throughput metric was obtained with 128 KB block size transfers. The throughput metric is useful in modeling large sequential transfers, for example, remote backup. The IOPs metric is useful in modeling small transactional traffic.

IOPs




Chart 1a shows the number of I/O operations per second for read and write operations on a single port and on both ports of an A6826A adapter measured on a 4 way rx4640 server. The X-axis represents the number of Fibre Channel ports and the Y-axis represents the number of I/O operations per second.

Chart 1a shows linear scaling of IOPs for one and two ports. The number of read I/O operations per second for a single port and for dual ports is 35180 and 70032 respectively. The number of write operations per second for a single port and for dual ports is 20619 and 41221 respectively. The IOPs metric is limited by the processor used on the A6826A adapter.

Chart 1b shows the %CPU utilization for I/O operation tests. The X-axis represents the number of Fibre Channel ports. The Y-axis represents the % CPU utilization. To accurately record CPU utilization, only one CPU was configured in these tests. Chart 1b demonstrates excellent CPU utilization numbers for small size I/Os on both ports of an A6826A adapter.

Service demand

Service Demand is a measure of the CPU time needed to handle 1 KB of I/O. The Service Demand number for an adapter is a measure of how much load the adapter places on the CPU(s) for each KB/s of throughput. The A6826A adapter offers low CPU service demand for small I/Os. To illustrate the service demand for the A6826A adapter, small size operations tests were conducted using db. Sequential Read and Write operations of 4KB and 8KB I/O sizes were performed on an rx4640 server with one CPU configured.

The following table shows single and dual port throughput and CPU utilization for raw Sequential Read and Sequential Write operations with 4KB and 8KB I/O sizes:

Number of Ports	Sequential Read						Sequential Write					
	4KB			8KB			4KB			8KB		
	Tput ¹	CPU ²	SD ³	Tput ¹	CPU ²	SD ³	Tput ¹	CPU ²	SD ³	Tput ¹	CPU ²	SD ³
Single	135	39%	2.8	189	30%	1.5	79	26%	3.2	111	21%	1.8
Dual	272	72%	2.6	379	55%	1.4	157	45%	2.8	223	36%	1.6

NOTE:

¹ Throughput in MB/s

² Single CPU utilization numbers

³ Service Demand in microseconds calculated using $((\%CPU \text{ utilization}/100)/ \text{Throughput in KB/s})(10^6)$

Data gathered on an rx4640server configured with a single CPU

Throughput

Chart 2a: Single and Dual Port Throughput





Chart 2b: CPU Utilization for Throughput Tests



Data gathered on an rx4640 server configured with a single CPU

Chart 2a shows the 128KB raw sequential read and write throughput of single and dual ports on an A6826A adapter. The X-axis represents the number of Fibre Channel ports. The Y-axis represents the throughput in MB/s.

Chart 2a shows one A6828A port performing read operations at 195MB/s. Two A6826A ports scale to 2x the performance of a single A6826A port, performing read operations at up to 390MB/s.

For write operations, one A6826A port performs at 186MB/s. Two A6826A ports scale to 2x the performance of a single A6826A port, performing writes at up to 372MB/s.


Chart 2a demonstrates outstanding read and write performance with linear scaling for 1 and 2 ports on the A6826A adapter.

Chart 2b shows the % CPU utilization in throughput tests. The X-axis represents the number of Fibre Channel ports. The Y-axis represents the % CPU utilization. To accurately record the CPU utilization, only one CPU is configured in these tests. Chart 2b shows very small percentages of CPU utilization for one or two A6826A ports in read and write throughput tests.

Scalability

A series of tests were conducted to evaluate the performance scalability of the A6826A adapter on rx2600, rx4640, rx7620, rx8620 and HP Integrity Superdome servers.

Chart 3a: 128KB Sequential Read Throughput




- rx2600 with two 900MHz Itanium 2 processors and 8GB Physical memory
- rx4640 with four 1.3GHz Itanium 2 processors and 8GB Physical memory
- rx7620, 2 cell partition, each cell with four 1.3 GHz Itanium 2 processors and 8GB Physical memory
- rx8620, 2 cell partition, each cell with four 1.3GHz Itanium 2 processors and 8GB Physical memory
- Integrity Superdome, 2 cell partition with eight 1.5GHz Itanium 2 processors and 8GB physical memory

Chart 3a shows the read throughput performance of 1, 2, 3 and 4 A6826A adapters installed in rx2600, rx4640, rx7620, rx8620, and HP Integrity Superdome servers. The X-axis represents the number of ports. The Y-axis represents the throughput performance in MB/s.

Chart 3a shows excellent read throughput performance with linear scaling for 4 A6826As installed in rx2600, rx4640, rx7620, rx8620 and HP Integrity Superdome servers, reaching an aggregate throughput of 1523MB/s. A6826A adapters installed in HP Integrity servers offer an excellent scalable SAN solution.

Chart 3b: 128 KB Sequential Write Throughput



- rx2600 with two 900MHz Itanium 2 processors and 8GB Physical memory
- rx4640 with four 1.3GHz Itanium 2 processors and 8GB Physical memory
- rx7620, 2 cell partition, each cell with four 1.3 GHz Itanium 2 processors and 8GB Physical memory
- rx8620, 2 cell partition, each cell with four 1.3GHz Itanium 2 processors and 8GB Physical memory
- Integrity Superdome, 2 cell partition with eight 1.5GHz Itanium 2 processors and 8GB physical memory

Chart 3b shows the write throughput performance of 1, 2, 3 and 4 A6826A adapters installed in rx2600, rx4640, rx7620, rx8620 and HP Integrity Superdome servers. The X-axis represents the number of ports. The Y-axis represents the throughput performance in MB/s.

Chart 3b shows excellent write throughput performance with linear scaling for 4 A6826A adapters installed in rx2600, rx4640, rx7620, rx8620 and HP Integrity Superdome servers, reaching an aggregate throughput of 1496 MB/s.

The table below summarizes the maximum number of cards that can be installed in the tested servers while maintaining linear scalability. Additional adapters may be installed to provide better connectivity.

Ports operating speed	Maximum number of A6826As for linear scalability				
	rx2600	rx4640	rx7620	rx8620	Integrity Superdome
Both Ports at 2Gb/s	4	4	2 per I/O Cage	2 per I/O Cage	2 per I/O Cage
One port at 1 Gb/s and other at 2Gb/s	4	4	3 per I/O Cage	3 per I/O Cage	3 per I/O Cage
Both ports at 1 Gb/s	4	6	4 per I/O Cage	4 per I/O Cage	4 per I/O Cage

System configuration guidelines

NOTE: Some PCI-X 133 MHz slots on Integrity Servers have higher bandwidth than other PCI-X 133 MHz slots on the server. In this whitepaper, we will refer to these higher bandwidth PCI-X 133 MHz slots as high performance PCI-X 133 MHz slots.

rx2600

The rx2600 is a 2-way entry-level Integrity server with 4 PCI-X 133MHz slots. A6826A adapters installed in an rx2600 server offer a great scalable solution.

rx4640

The rx4640 is a 4 way entry level Integrity server with eight I/O slots. Slots 1 and 2 are reserved for core SCSI and core LAN. Slots 3, 4, 5, and 6 are shared slots operating at 3.3v, 64-bit PCI-66MHz. Slots 7 and 8 are high performance PCI-X 133MHz slots. In order to achieve linear scalability with up to four A6826A adapters, HP recommends the first two A6826As be installed in slot 7 and slot 8, the next two A6826As should be installed in the shared slots, but avoid installing A6826A adapters in adjacent shared slots.

rx7620 / rx8620

The rx7620 and rx8620 are cell based mid-range Integrity servers. The I/O subsystem in the rx7620 and rx8620 offers 16 PCI-X 133MHz slots when these servers are configured with 2 cells. Fourteen of the sixteen slots are high performance PCI-X 133MHz slots. HP recommends A6826A adapters be installed in high performance PCI-X 133MHz slots (Physical slot numbers 1, 2, 3, 4, 5, 6, 7) to achieve the performance demonstrated in this paper. To obtain linear scalability, HP recommends the rx7620 or rx8620 be configured with a 2-cell partition: Two A6826As should be installed in any of the PCI-X slots numbered 1 to 7 and a maximum of two A6826As should be installed in any of the PCI-X slots numbered 9 to 15. Each cell should be configured with a multiple of 8 equal capacity DIMMs to take advantage of memory interleaving. The DIMMs on a cell should be evenly distributed across the busses.







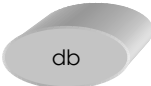

HP Integrity Superdome

The HP Integrity Superdome is a cell based high-end Integrity server. Each Superdome I/O Cage subsystem offers 4 high performance PCI-X slots capable of delivering 1066MB/s peak bandwidth. HP recommends A6826A adapters be installed in high performance PCI-X slots (Physical slot numbers 4, 5, 6, 7) to achieve the performance demonstrated in this paper. To obtain the linear scalability demonstrated in this paper, an Integrity Superdome partition should be configured with at least two cells with 8 Gb of physical memory and not more than two A6826As should be installed in any of the four high performance slots in each I/O cage. Each cell should be configured with a multiple of 8 equal capacity DIMMs to take advantage of memory interleaving. The DIMMs on a cell should be evenly distributed across the busses.

Test details

The performance results presented in this paper were obtained with A6826A adapters installed in various HP Integrity servers. The system configurations used in the test setup are detailed in the following table:

Products used in testing


products used for performance measurement test		
Servers Tested		rx2600 <ul style="list-style-type: none"> Configured with <ul style="list-style-type: none"> Two 900MHz Intel Itanium 2 Processors 4GB System Memory HP-UX 11i v2 OS A6826A Fibre Channel driver (Ver. B.11.23.01)
		rx4640 <ul style="list-style-type: none"> Configured with <ul style="list-style-type: none"> Four 1.3GHz Intel Itanium 2 Processors 8 GB System Memory HP-UX 11i v2 OS A6826A Fibre Channel driver (Ver. B.11.23.01)
		rx7620 <ul style="list-style-type: none"> 2 Cell partition, each cell with <ul style="list-style-type: none"> Four 1.3GHz Intel Itanium 2 Processors 8GB System Memory HP-UX 11i v2 OS A6826A Fibre Channel driver (Ver. B.11.23.01)
		rx8620 <ul style="list-style-type: none"> 2 Cell partition, each cell with <ul style="list-style-type: none"> Four 1.3GHz Intel Itanium 2 Processors 8GB System Memory HP-UX 11i v2 OS A6826A Fibre Channel driver (Ver. B.11.23.01)
		Integrity Superdome <ul style="list-style-type: none"> 2 Cell partition, each cell with <ul style="list-style-type: none"> Four 1.5GHz Intel Itanium 2 Processors 8GB System Memory HP-UX 11i v2 OS A6826A Fibre Channel driver (Ver. B.11.23.01)
		A6826A HBA <ul style="list-style-type: none"> PCI-X dual-channel 2Gb Fibre Channel adapter Each port capable of independently operating at 1 or 2 Gb/s (Auto-Negotiation) 33/66/100/133MHz-64bit capable
Benchmark software		Diskbench (db) is the benchmark suite that generated disk read and write traffic for these tests.
DS2405		HP StorageWorks 2Gb/s Disk System

Test configuration

The rx4640 test configuration consisted of a 4-way rx4640 with four 1.3GHz processors, 8 GB of system memory and 4 A6826A adapters installed. Two A6826A adapters were connected to one HP 2Gb/s Fibre Channel switch and the other two A6826A adapters were connected to the another HP 2Gb/s Fibre Channel switch. Sixteen HP DS2405s were evenly connected to the 2 switches. The switches were zoned so each zone consisted of a single A6828A port and 2 DS2405s. A Single CPU was configured on the rx4640 to obtain the CPU utilization demonstrated in the IOPs and throughput tests for single and dual A6828A ports. Additional CPUs were added to study the scalability of the A6826A on the rx4640.

A similar test setup was used to collect performance data on the rx2600, rx7620, rx8620 and Integrity Superdome.

test setup



Additional information

For more information about the A6826A adapter please visit:
<http://www.hp.com/products1/serverconnectivity/storagesnf2/index.html>.

For more information about HP's Itanium-based Integrity servers please visit:
<http://www.hp.com/products1/servers/integrity/index.html>

© 2005 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries and are used under license. UNIX is a registered trademark of The Open Group.

02/2005

